

SAMPLE-PATH LARGE DEVIATIONS IN CREDIT RISK

V. LEIJDEKKER, M. MANDJES, AND P. SPREIJ

ABSTRACT. The event of large losses plays an important role in credit risk. As these large losses are typically rare, and portfolios usually consist of a large number of positions, large deviation theory is the natural tool to analyze the tail asymptotics of the probabilities involved. We first derive a sample-path large deviation principle (LDP) for the portfolio's loss process, which enables the computation of the logarithmic decay rate of the probabilities of interest. In addition, we derive exact asymptotic results for a number of specific rare-event probabilities, such as the probability of the loss process exceeding some given function.

1. INTRODUCTION

For financial institutions, such as banks and insurance companies, it is of crucial importance to accurately assess the risk of their portfolios. These portfolios typically consist of a large number of obligors, such as mortgages, loans or insurance policies, and therefore it is computationally infeasible to treat each individual object in the portfolio separately. As a result, attention has shifted to measures that characterize the risk of the portfolio as a whole, see e.g. [DS03] for general principles concerning managing credit risk. The best known metric is the so-called *value at risk*, see [Jor06], measuring the minimum amount of money that can be lost with α percent certainty over some given period. Several other measures have been proposed, such as economic capital, the risk adjusted return on capital (RAROC), or expected shortfall, which is a coherent risk measure [ADEH99]. Each of these measure are applicable to market risk as well as credit risk. Measures such as loss given default (LGD) and exposure at default (EAD) are measures that purely apply to credit risk. These and other measures are discussed in detail in e.g. [MFE05].

The currently existing methods mainly focus on the distribution of the portfolio loss *up to a given point in time* (for instance one year into the future). It can be argued, however, that in many situations it makes more sense to use probabilities that involve the (cumulative) loss *process*, say $\{L(t) : t \geq 0\}$. Highly relevant, for instance, is the event that $L(\cdot)$ ever exceeds a given function $\zeta(\cdot)$ (within a certain time window, for instance between now and one year ahead), i.e., an event of the type

$$(1.1) \quad \{\exists t \leq T : L(t) \geq \zeta(t)\}.$$

It is clear that measures of the latter type are intrinsically harder to analyze, as it does not suffice anymore to have knowledge of the marginal distribution of the loss process at a given point in time; for instance the event (1.1) actually corresponds to the union of events $\{L(t) \geq \zeta(t)\}$, for $t \leq T$, and its probability will depend on the law of $L(\cdot)$ as a process on $[0, T]$.

In line with the remarks we made above, earlier papers on applications of large-deviation theory to credit risk, mainly address the (asymptotics of the) distribution of the loss process at a single point in time, see e.g. [DDD04] and [GKS07]. The former paper considers, in addition, also

the probability that the *increments* of the loss process exceed a certain level. Other approaches to quantifying the tail distribution of the losses have been taken by [LKSS01], who use extreme-value theory (see [EKM97] for a background), [Gor02] and [MTB01], where the authors consider saddle point approximations to the tails of the loss distribution. Numerical and simulation techniques for credit risk can be found in e.g. [Gla04]. The first contribution of our work concerns a so-called sample-path large deviation principle (LDP) for the average cumulative losses for large portfolios. Loosely speaking, such an LDP means that, with $L_n(\cdot)$ denoting the loss process when n obligors are involved, we can compute the logarithmic asymptotics (for n large) of the normalized loss process $L_n(\cdot)/n$ being in a set of trajectories A :

$$(1.2) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right);$$

we could for instance pick a set A that corresponds to the event (1.1). Most of the sample-path LDPs that have been developed so far involve stochastic processes with independent or nearly-independent increments, see for instance the results by Mogul'skiĭ for random walks [Mog76], de Acosta for Lévy processes [Aco94], and Chang [Cha95] for weakly-correlated processes; results for processes with a stronger correlation structure are restricted to special classes of processes, such as Gaussian processes, see e.g. [Aze80]. It is observed that our loss process is not covered by these results, and therefore new theory had to be developed. The proof of our LDP relies on 'classical' large-deviation results (such as Cramér's theorem, Sanov's theorem, Mogul'skiĭ's theorem), but in addition the concept of *epi-convergence* [Kal86] is relied upon.

Our second main result focuses specifically on the event (1.1) of ever (before some time horizon T) exceeding a given barrier function $\zeta(\cdot)$. Whereas we so far considered, inherently imprecise, logarithmic asymptotics of the type displayed in (1.2), we can now compute so-called *exact asymptotics*: we identify an explicit function $f(n)$ such that $f(n)/p_n \rightarrow 1$ as $n \rightarrow \infty$, where p_n is the probability of our interest. As is known from the literature, it is in general substantially harder to find exact asymptotics than logarithmic asymptotics. The proof uses the fact that, after discretizing time, the contribution of just a single time epoch dominates, in the sense that there is a t^* such that

$$(1.3) \quad \mathbb{P} \left(\frac{1}{n} L_n(t^*) \geq \zeta(t^*) \right) / p_n \rightarrow 1, \quad \text{with } p_n := \mathbb{P} \left(\exists t : \frac{1}{n} L_n(t) \geq \zeta(t) \right).$$

This t^* can be interpreted as the most likely epoch of exceeding $\zeta(\cdot)$.

Turning back to the setting of credit risk, both of the results we present are derived in a setup where all obligors in the portfolio are i.i.d., in the sense that they behave independently and stochastically identically. A third contribution of our work concerns a discussion on how to extend our results to cases where the obligors are *dependent* (meaning that they, in the terminology of [DDD04], react to the same 'macro-environmental' variable, conditional upon which they are independent again). We also treat the case of obligor-heterogeneity: we show how to extend the results to the situation of multiple classes of obligors.

The paper is structured as follows. In Section 2 we introduce the loss process and we describe the scaling under which we work. We also recapitulate a couple of relevant large-deviation results. Our first main result, the sample-path LDP for the cumulative loss process, is stated and proved in Section 3. Special attention is paid to, easily-checkable, sufficient conditions under which this result holds. As argued above, the LDP is a generally applicable result, as it yields an expression for the decay rate of any probability that depends on the entire sample path. Then, in Section 4, we derive the exact asymptotic behavior of the probability that, at some point in time, the loss

exceeds a certain threshold, i.e., the asymptotics of p_n , as defined in (1.3). After this we derive a similar result for the *increments* of the loss process. Eventually, in Section 5, we discuss a number of possible extensions to the results we have presented. Special attention is given to allowing dependence between obligors, and to different classes of obligors each having its own specific distributional properties.

2. NOTATION AND DEFINITIONS

The portfolios of banks and insurance companies are typically very large: they may consist of several thousands of assets. It is therefore computationally impossible to estimate the risks for each element, or obligor, in a portfolio. This explains why one attempts to assess the aggregated losses resulting from defaults, e.g. bankruptcies, failure to repay loans or insurance claims, for the portfolio as a whole. The risk in the portfolio is then measured through this (aggregate) loss process. In the following sections we introduce the loss process and the portfolio constituents more formally.

2.1. Loss Process. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability space on which all random variables below are defined. We assume that the portfolio consists of n obligors and we denote the default time of obligor i by τ_i . Further we write U_i for the loss incurred on a default of obligor i . We then define the cumulative loss process L_n as

$$(2.1) \quad L_n(t) := \sum_{i=1}^n U_i Z_i(t),$$

where $Z_i(t) = 1_{\{\tau_i \leq t\}}$ is the default indicator of obligor i . We assume that the loss amounts $U_i \geq 0$ are i.i.d., and that the default times $\tau_i \geq 0$ are i.i.d. as well. In addition we assume that the loss amounts and the default times are mutually independent. In the remainder of this paper U and $Z(t)$ denote generic random variables with the same distribution as the U_i and $Z_i(t)$, respectively.

Throughout this paper we assume that the defaults only occur on the time grid \mathbb{N} ; in Section 5 we discuss how to deal with the default epochs taking continuous values. In some cases we explicitly consider a finite time grid, say $\{1, 2, \dots, N\}$. The extension of the results we derive to a more general grid $\{0 < t_1 < t_2 < \dots < t_N\}$ is completely trivial. The distribution of the default times, for each j , is denoted by

$$(2.2) \quad p_j := \mathbb{P}(\tau = j),$$

$$(2.3) \quad F_j := \mathbb{P}(\tau \leq j) = \sum_{i=1}^j p_i.$$

Given the distribution of the loss amounts U_i and the default times τ_i , our goal is to investigate the loss process. Many of the techniques that have been developed so far, first fix a time T (typically one year), and then stochastic properties of the cumulative loss at time T , i.e., $L_n(T)$, are studied. Measures such as value at risk and economic capital are examples of these ‘one-dimensional’ characteristics. Many interesting measures, however, involve properties of the entire *path* of the loss process rather than those of just one time epoch, examples being the probability that that $L_n(\cdot)$ exceeds some barrier function $\zeta(\cdot)$ for some t smaller than the horizon T , or the probability that (during a certain period) the loss always stays above a certain level. The event corresponding to the former probability might require the bank to attract more capital, or worse, it might lead to the bankruptcy of this bank. The event corresponding to the latter event might also lead to the

bankruptcy of the bank, as a long period of stress may have substantial negative implications. We conclude that having a handle on these probabilities is therefore a useful instrument when assessing the risks involved in the bank's portfolios.

As mentioned above, the number of obligors n in a portfolio is typically very large, thus prohibiting analyses based on the specific properties of the individual obligors. Instead, it is more natural to study the asymptotical behavior of the loss process as $n \rightarrow \infty$. One could rely on a central-limit-theorem based approach, but in this paper we focus on rare events, by using the theory of large deviations.

In the following subsection we provide some background of large-deviation theory, and we define a number of quantities that are used in the remainder of this paper.

2.2. Large Deviation Principle. In this section we give a short introduction to the theory of large deviations. Here, in an abstract setting, the limiting behavior of a family of probability measures $\{\mu_n\}$ on the Borel sets \mathcal{B} of a metric space (\mathcal{X}, d) is studied, as $n \rightarrow \infty$. This behavior is referred to as the Large Deviation Principle (LDP), and it is characterized in terms of a *rate function*. The LDP states lower and upper exponential bounds for the value that the measures μ_n assign to sets in a topological space \mathcal{X} . Below we state the definition of the rate function that has been taken from [DZ98].

Definition 2.1. *A rate function is a lower semicontinuous mapping $I : \mathcal{X} \rightarrow [0, \infty]$, for all $\alpha \in [0, \infty)$ the level set $\Psi_I(\alpha) := \{x \mid I(x) \leq \alpha\}$ is a closed subset of \mathcal{X} . A good rate function is a rate function for which all the level sets are compact subsets of \mathcal{X} .*

With the definition of the rate function in mind we state the large deviation principle for the sequence of measure $\{\mu_n\}$.

Definition 2.2. *We say that $\{\mu_n\}$ satisfies the large deviation principle with a rate function $I(\cdot)$ if*

(i) *(Upper bound) for any closed set $F \subseteq \mathcal{X}$*

$$(2.4) \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mu_n(F) \leq - \inf_{x \in F} I(x).$$

(ii) *(Lower bound) for any open set $G \subseteq \mathcal{X}$*

$$(2.5) \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mu_n(G) \geq - \inf_{x \in G} I(x).$$

We say that a family of random variables $X = \{X_n\}$, with values in \mathcal{X} , satisfies an LDP with rate function $I_X(\cdot)$ iff the laws $\{\mu_n^X\}$ satisfy an LDP with rate function I_X , where μ_n^X is the law of X_n .

The so-called Fenchel-Legendre transform plays an important role in expressions for the rate function. Let for an arbitrary random variable X , the logarithmic moment generating function, sometimes referred to as cumulant generating function, be given by

$$(2.6) \quad \Lambda_X(\theta) := \log M_X(\theta) = \log \mathbb{E} [e^{\theta X}] \leq \infty,$$

for $\theta \in \mathbb{R}$. The Fenchel-Legendre transform Λ_X^* of Λ_X is then defined by

$$(2.7) \quad \Lambda_X^*(x) := \sup_{\theta} (\theta x - \Lambda_X(\theta)).$$

We sometimes say that Λ_X^* is the Fenchel-Legendre transform of X .

The LDP from Definition 2.2 provides upper and lower bounds for the log-asymptotic behavior of measures μ_n . In case of the loss process (2.1), *fixed* at some time t , we can easily establish an LDP by an application of Cramér's theorem (Theorem A.1). This theorem yields that the rate function is given by $\Lambda_{UZ(t)}^*(\cdot)$, where $\Lambda_{UZ(t)}^*(\cdot)$ is the Fenchel-Legendre transform of the random variable $UZ(t)$.

The results we present in this paper involve either $\Lambda_U^*(\cdot)$ (Section 3), which corresponds to i.i.d. loss amounts U_i only, or $\Lambda_{UZ(t)}^*(\cdot)$ (Section 4), which corresponds to those loss amounts up to time t . In the following section we derive an LDP for the whole path of the loss process, which can be considered as an extension of Cramér's theorem.

3. A SAMPLE-PATH LARGE DEVIATION RESULT

In the previous section we have introduced the large deviation principle. In this section we derive a sample-path LDP for the cumulative loss process (2.1). We consider the exponential decay of the probability that the path of the loss process $L_n(\cdot)$ is in some set A , as the size n of the portfolio tends to infinity.

3.1. Assumptions. In order to state a sample-path LDP we need to define the topology that we work on. To this end we define the space \mathcal{S} of all nonnegative and nondecreasing functions on $T_N = \{1, 2, \dots, N\}$,

$$\mathcal{S} := \{f : T_N \rightarrow \mathbb{R}_0^+ \mid 0 \leq f_i \leq f_{i+1} \text{ for } i < N\}.$$

This set is identified with the space $\mathbb{R}_{\leq}^N := \{x \in \mathbb{R}^N \mid 0 \leq x_i \leq x_{i+1} \text{ for } i < N\}$. The topology on this space is the one induced by the supremum norm

$$\|f\|_{\infty} = \max_{i=1, \dots, N} |f_i|.$$

As we work on a finite-dimensional space, the choice of the norm is not important, as any other norm on \mathcal{S} would result in the same topology. We use the supremum norm as this is convenient in some of the proofs in this section.

We identify the space of all probability measures on T_N with the simplex Φ :

$$(3.1) \quad \Phi := \left\{ \varphi \in \mathbb{R}^N \mid \sum_{i=1}^N \varphi_i = 1, \varphi_i \geq 0 \text{ for } i \leq N \right\}.$$

For a given $\varphi \in \Phi$ we denote the cumulative distribution function by ψ , i.e.,

$$(3.2) \quad \psi_i = \sum_{j=1}^i \varphi_j, \text{ for } i \leq N;$$

note that $\psi \in \mathcal{S}$ and $\psi_N = 1$.

Furthermore, we consider the loss amounts U_i as introduced in Section 2.1, a $\varphi \in \Phi$ with cdf ψ , and a sequence of $\varphi^n \in \Phi$, each with cdf ψ^n , such that $\varphi^n \rightarrow \varphi$ as $n \rightarrow \infty$, meaning that $\varphi_i^n \rightarrow \varphi_i$

for all $i \leq N$. We define two families of measures (μ_n) and (ν_n) ,

$$(3.3) \quad \mu_n(A) := \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{\lfloor n\psi_i \rfloor} U_j \right)_{i=1}^N \in A \right),$$

$$(3.4) \quad \nu_n(A) := \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{\lfloor n\psi_i^n \rfloor} U_j \right)_{i=1}^N \in A \right),$$

where $A \in \mathcal{B} := \mathcal{B}(\mathbb{R}^N)$ and $[x] := \sup \{k \in N \mid k \leq x\}$. Below we state an assumption under which the main result in this section holds. This assumption refers to the definition of *exponential equivalence*, which can be found in Definition A.2.

Assumption 3.1. *Assume that $\varphi^n \rightarrow \varphi$ and moreover that the measures μ_n and ν_n as defined in (3.3) and (3.4), respectively, are exponentially equivalent.*

From Assumption 3.1 we learn that the differences between the two measures μ_n and ν_n go to zero at a ‘superexponential’ rate. In the next section, in Lemma 3.4, we provide a sufficient condition, that is easy to check, under which this assumption holds.

3.2. Main Result. The assumptions and definitions in the previous sections allow us to state the main result of this section. We show that the average loss process satisfies a large deviation principle as in (2.2). This principle allows us to approximate a large variety of probabilities related to the average loss process, such as the probability that the loss process stays above a certain time dependent level or the probability that the loss process exceeds a certain level before some given point in time.

Theorem 3.2. *With Φ as in (3.1) and under Assumption 3.1, the average loss process, $L_n(\cdot)/n$ satisfies an LDP with rate function $I_{U,p}$. Here, for $x \in \mathbb{R}_{\leq}^N$, $I_{U,p}$ is given by*

$$(3.5) \quad I_{U,p}(x) := \inf_{\varphi \in \Phi} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right),$$

with $\Delta x_i := x_i - x_{i-1}$ and $x_0 := 0$.

Observing the rate function for this sample path LDP, we see that the effects of the default times τ_i and the loss amounts U_i are nicely decomposed into the two terms in the rate function, one involving the distribution of the default epoch τ (the ‘Sanov term’, cf. [DZ98, Thm. 6.2.10]), the other one involving the incurred loss size U (the ‘Cramér term’, cf. [DZ98, Thm. 2.2.3]). Observe that we recover Cramér’s theorem by considering a time grid *consisting of a single time point*, which means that Theorem 3.2 extends Cramér’s result. We also remark that, informally speaking, the optimizing $\varphi \in \Phi$ in (3.5) can be interpreted as the ‘most likely’ distribution of the loss epoch, given that the path of $L_n(\cdot)/n$ is close to x .

As a sanity check we calculate the value of the rate function $I_{U,p}(x)$ for the ‘average path’ of $L_n(\cdot)/n$, given by $x_j^* = \mathbb{E}[U] F_j$ for $j \leq N$, where F_j is the cumulative distribution of the default times as given in (2.3); this path should give a rate function equal to 0. To see this we first remark that clearly $I_{U,p}(x) \geq 0$ for all x , since both the Sanov term and the Cramér term are non-negative.

This yields the following chain of inequalities:

$$\begin{aligned}
0 &\leq I_{U,p}(x^*) = \inf_{\varphi \in \Phi} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\mathbb{E}[U] p_i}{\varphi_i} \right) \right) \\
&\leq \sum_{i=1}^N p_i \left(\log \left(\frac{p_i}{p_i} \right) + \Lambda_U^* \left(\frac{\mathbb{E}[U] p_i}{p_i} \right) \right) \\
&= \sum_{i=1}^N p_i \Lambda_U^* (\mathbb{E}[U]) = \Lambda_U^* (\mathbb{E}[U]) = 0,
\end{aligned}$$

where we have used that for $\mathbb{E}[U] < \infty$, it always holds that $\Lambda_U^* (\mathbb{E}[U]) = 0$ cf. [DZ98, Lemma 2.2.5]. The inequalities above thus show that, if the ‘average path’ x^* lies in the set of interest, then the corresponding decay rate is 0, meaning that the probability of interest decays subexponentially.

In the proof of Theorem 3.2 we use the following lemma, which is related to the concept of *epi-convergence*, extensively discussed in [Kal86].

Lemma 3.3. *Let $f_n, f : D \rightarrow \mathbb{R}$, with $D \subset \mathbb{R}^m$ compact. Assume that for all $x \in D$ and for all $x_n \rightarrow x$ in D we have*

$$(3.6) \quad \limsup_{n \rightarrow \infty} f_n(x_n) \leq f(x).$$

Then we have

$$\limsup_{n \rightarrow \infty} \sup_{x \in D} f_n(x) \leq \sup_{x \in D} f(x).$$

Proof. Let $f_n^* = \sup_{x \in D} f_n(x)$, $f^* = \sup_{x \in D} f(x)$. Consider a subsequence $f_{n_k}^* \rightarrow \limsup_{n \rightarrow \infty} f_n^*$. Let $\varepsilon > 0$ and choose x_{n_k} such that $f_{n_k}^* < f_{n_k}(x_{n_k}) + \varepsilon$ for all k . By the compactness of D , there exists a limit point x such that along a subsequence $x_{n_{k_j}} \rightarrow x$. By the hypothesis (3.6) we then have

$$\limsup_{n \rightarrow \infty} f_n^* \leq \lim_{n \rightarrow \infty} f_{n_{k_j}}(x_{n_{k_j}}) + \varepsilon = f(x) + \varepsilon \leq f^* + \varepsilon.$$

Let $\varepsilon \downarrow 0$ to obtain the result. \square

Proof of Theorem 3.2. We start by establishing an identity from which we show both bounds. We need to calculate the probability

$$\mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right) = \mathbb{P} \left(\left(\frac{1}{n} L_n(1), \dots, \frac{1}{n} L_n(N) \right) \in A \right),$$

for certain $A \in \mathcal{B}$. For each point j on the time grid T_N we record by the ‘default counter’ $K_{n,j} \in \{0, \dots, n\}$ the number of defaults at time j :

$$K_{n,j} := \#\{i \in \{1, \dots, n\} \mid \tau_i = j\}.$$

These counters allow us to rewrite the probability to

$$\begin{aligned}
(3.7) \quad \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right) &= \mathbb{E} \left[\mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{K_{n,1}} U_{(j)}, \dots, \frac{1}{n} \sum_{j=1}^{K_{n,1} + \dots + K_{n,N}} U_{(j)} \right) \in A \mid K_n \right) \right] \\
&= \sum_{k_1 + \dots + k_N = n} \mathbb{P}(K_{n,i} = k_i \text{ for } i \leq N) \times \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{m_i} U_{(j)} \right)_{i=1}^N \in A \right),
\end{aligned}$$

where $m_i := \sum_{j=1}^i k_j$ and the loss amounts U_j have been ordered, such that the first $U_{(j)}$ correspond to the losses at time 1, etc.

Upper bound. Starting from Equality (3.7), let us first establish the upper bound of the LDP. To this end, let F be a closed set and consider the decay rate

$$(3.8) \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in F \right).$$

An application of Lemma A.3 together with (3.7), implies that (3.8) equals

$$(3.9) \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in F \right) = \limsup_{n \rightarrow \infty} \max_{\sum k_j = n} \frac{1}{n} \left[\log \mathbb{P} \left(\frac{K_{n,i}}{n} = \frac{k_i}{n}, i \leq N \right) + \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{m_i} U_{(j)} \right)_{i=1}^N \in F \right) \right],$$

Next we replace the dependence on n in the maximization by maximizing over the set Φ as in (3.1). In addition, we replace the k_i in (3.9) by

$$(3.10) \quad \hat{\varphi}_{n,i} := \frac{[n\psi_i] - [n\psi_{i-1}]}{n},$$

where the ψ_i have been defined in (3.2). As a result, (3.8) reads

$$(3.11) \quad \limsup_{n \rightarrow \infty} \sup_{\varphi \in \Phi} \frac{1}{n} \left[\log \mathbb{P} \left(\frac{K_{n,i}}{n} = \hat{\varphi}_{n,i}, i \leq N \right) + \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in F \right) \right].$$

Note that (3.8) equals (3.11), since for each n and vector $(k_1, \dots, k_N) \in \mathbb{N}^N$, with $\sum_{i=1}^N k_i = n$, there is a $\varphi \in \Phi$ with $\varphi_i = k_i/n$. On the other hand, we only cover outcomes of this form by rounding off the φ_i .

We can bound the first term in this expression from above using Lemma A.5, which implies that the decay rate (3.8) is majorized by

$$\limsup_{n \rightarrow \infty} \sup_{\varphi \in \Phi} \left[- \sum_{i=1}^N \hat{\varphi}_{n,i} \log \left(\frac{\hat{\varphi}_{n,i}}{p_i} \right) + \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in F \right) \right].$$

Now note that calculating the limsup in the previous expression is not straightforward due to the supremum over Φ . The idea is therefore to interchange the supremum and the limsup, by using Lemma 3.3. To apply this lemma we first introduce

$$f_n(\varphi) := - \sum_{i=1}^N \varphi_i \log \left(\frac{\varphi_i}{p_i} \right) + \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in F \right),$$

$$f(\varphi) := - \inf_{x \in F} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right),$$

and note that Φ is a compact subset of \mathbb{R}^n . We have to show that for any sequence $\varphi^n \rightarrow \varphi$

$$(3.12) \quad \limsup_{n \rightarrow \infty} f_n(\varphi^n) \leq f(\varphi),$$

such that the conditions of Lemma 3.3 are satisfied. We observe, with ψ_i^n as in (3.2) and $\hat{\varphi}_i^n$ as in (3.10) with φ replaced by φ^n , that

$$\begin{aligned} \limsup_{n \rightarrow \infty} f_n(\varphi^n) &\leq \limsup_{n \rightarrow \infty} \left(- \sum_{i=1}^N \hat{\varphi}_{n,i}^n \log \left(\frac{\hat{\varphi}_{n,i}^n}{p_i} \right) \right) \\ &\quad + \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{\lfloor n\psi_i^n \rfloor} U_{(j)} \right)_{i=1}^N \in F \right). \end{aligned}$$

Since $\varphi^n \rightarrow \varphi$ and since $\hat{\varphi}_{n,i}^n$ differs at most by $1/n$ from φ_i^n , it immediately follows that $\hat{\varphi}_{n,i}^n \rightarrow \varphi_i$. For an arbitrary continuous function g we thus have $g(\hat{\varphi}_{n,i}^n) \rightarrow g(\varphi_i)$. This implies that

$$(3.13) \quad \limsup_{n \rightarrow \infty} \left(- \sum_{i=1}^N \hat{\varphi}_{n,i}^n \log \left(\frac{\hat{\varphi}_{n,i}^n}{p_i} \right) \right) = - \sum_{i=1}^N \varphi_i \log \left(\frac{\varphi_i}{p_i} \right).$$

Inequality (3.12) is established once we have shown that

$$(3.14) \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{\lfloor n\psi_i^n \rfloor} U_{(j)} \right)_{i=1}^N \in F \right) \leq - \inf_{x \in F} \sum_{i=1}^N \varphi_i \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right).$$

By Assumption 3.1, we can exploit the exponential equivalence together with Theorem A.7, to see that (3.14) holds as soon as we have that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{\lfloor n\psi_i \rfloor} U_{(j)} \right)_{i=1}^N \in F \right) \leq - \inf_{x \in F} \sum_{i=1}^N \varphi_i \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right).$$

But this inequality is a direct consequence of Lemma A.6, and we conclude that (3.14) holds. Combining (3.13) with (3.14) yields

$$\begin{aligned} \limsup_{n \rightarrow \infty} f_n(\varphi^n) &\leq - \sum_{i=1}^N \varphi_i \log \left(\frac{\varphi_i}{p_i} \right) - \inf_{x \in F} \sum_{i=1}^N \varphi_i \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \\ &= - \inf_{x \in F} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right) = f(\varphi), \end{aligned}$$

so that indeed the conditions of Lemma 3.3 are satisfied, and therefore

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{\varphi \in \Phi} f_n(\varphi) &\leq \sup_{\varphi \in \Phi} f(\varphi) = \sup_{\varphi \in \Phi} \left(- \inf_{x \in F} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right) \right) \\ &= - \inf_{x \in F} \inf_{\varphi \in \Phi} \sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right) = - \inf_{x \in F} I_{U,p}(x). \end{aligned}$$

This establishes the upper bound of the LDP.

Lower bound. To complete the proof, we need to establish the corresponding lower bound. Let G be an open set and consider

$$(3.15) \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in G \right).$$

We apply Equality (3.7) to this liminf, with A replaced by G , and we observe that this sum is larger than the largest term in the sum, which shows that (where we directly switch to the enlarged

space Φ) the decay rate (3.15) majorizes

$$\liminf_{n \rightarrow \infty} \sup_{\varphi \in \Phi} \frac{1}{n} \left(\log \mathbb{P} \left(\frac{1}{n} K_{n,i} = \hat{\varphi}_{n,i}, i \leq N \right) + \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in G \right) \right).$$

Observe that for any sequence of functions $h_n(\cdot)$ it holds that $\liminf_n \sup_x h_n(x) \geq \liminf_n h_n(\tilde{x})$ for all \tilde{x} , so that we obtain the evident inequality

$$\liminf_{n \rightarrow \infty} \sup_x h_n(x) \geq \sup_x \liminf_{n \rightarrow \infty} h_n(x).$$

This observation yields that the decay rate of interest (3.15) is not smaller than

$$(3.16) \quad \sup_{\varphi \in \Phi} \left(\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} K_{n,i} = \hat{\varphi}_{n,i}, i \leq N \right) + \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in G \right) \right),$$

where we have used that $\liminf_n (x_n + y_n) \geq \liminf_n x_n + \liminf_n y_n$. We apply Lemma A.5 to the first liminf in (3.16), leading to

$$(3.17) \quad \begin{aligned} & \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} K_{n,i} = \hat{\varphi}_{n,i}, i \leq N \right) \geq \liminf_{n \rightarrow \infty} \left(- \sum_{i=1}^N \hat{\varphi}_{n,i} \log \left(\frac{\hat{\varphi}_{n,i}}{p_i} \right) - \frac{N}{n} \log(n+1) \right) \\ & = - \sum_{i=1}^N \varphi_i \log \left(\frac{\varphi_i}{p_i} \right), \end{aligned}$$

since $\log(n+1)/n \rightarrow 0$ as $n \rightarrow \infty$. The second liminf in (3.16) can be bounded from below by an application of Lemma A.6. Since G is an open set, this lemma yields

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_{(j)} \right)_{i=1}^N \in G \right) \geq - \inf_{x \in G} \sum_{i=1}^N \varphi_i \Lambda_U^* \left(\frac{x_i - x_{i-1}}{\varphi_i} \right).$$

Upon combining these two results, we see that we have established the lower bound:

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in G \right) \\ & \geq - \inf_{\varphi \in \Phi} \inf_{x \in A} \left(\sum_{i=1}^N \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{x_i - x_{i-1}}{\varphi_i} \right) \right) \right) = - \inf_{x \in G} I_{U,p}(x). \end{aligned}$$

This completes the proof of the theorem. \square

In order to apply Theorem 3.2, one needs to check that Assumption 3.1 holds. In general this could be a quite cumbersome exercise. In Lemma 3.4 below we provide a sufficient, easy-to-check condition under which this assumption holds.

Lemma 3.4. *Assume that for all $\theta \in \mathbb{R} : \Lambda_U(\theta) < \infty$. Then Assumption 3.1 holds.*

Remark 3.5. The assumption we make in Lemma 3.4, i.e., that the logarithmic moment generating function is finite everywhere, is a common assumption in large deviations theory. We remark that for instance Mogul'skiĭ's theorem [DZ98, Thm. 5.1.2] also relies on this assumption; this theorem is a sample-path LDP for

$$Y_n(t) := \frac{1}{n} \sum_{i=1}^{[nt]} X_i,$$

on the interval $[0, 1]$. In Mogul'skiĭ's result the X_i are assumed to be i.i.d; in our model we have that $L_n(t) = \sum_{i=1}^n U_i Z_i(t)/n$, so that our sample-path result clearly does not fit into the setup of Mogul'skiĭ's theorem. \diamond

Remark 3.6. In Lemma 3.4 it was assumed that $\Lambda_U(\theta) < \infty$, for all $\theta \in \mathbb{R}$, but an equivalent condition is

$$(3.18) \quad \lim_{x \rightarrow \infty} \frac{\Lambda_U^*(x)}{x} = \infty.$$

In other words: this alternative condition can be used instead of the condition stated in Lemma 3.4. To see that both requirements are equivalent, make the following observations. In Lemma A.4 it is shown that (3.18) is implied by the assumption in Lemma 3.4. In order to prove the converse, assume that (3.18) holds, and that there is a $0 < \theta_0 < \infty$ for which $\Lambda_U(\theta) = \infty$. Without loss of generality we can assume that $\Lambda_U(\theta)$ is finite for $\theta < \theta_0$ and infinite for $\theta \geq \theta_0$. For $x > \mathbb{E}[U]$, the Fenchel-Legendre transform is then given by

$$\Lambda_U^*(x) = \sup_{0 < \theta < \theta_0} (\theta x - \Lambda_U(\theta)).$$

Since $U \geq 0$ and $\Lambda_U(0) = 0$, we know that $\Lambda_U(\theta) \geq 0$ for $0 < \theta < \theta_0$, and hence

$$\frac{\Lambda_U^*(x)}{x} \leq \theta_0,$$

which contradicts the assumption that this ratio tends to infinity as $x \rightarrow \infty$, and thus establishing the equivalence. \diamond

Proof of Lemma 3.4. Let $\varphi^n \rightarrow \varphi$ for some sequence of $\varphi^n \in \Phi$ and $\varphi \in \Phi$. We introduce two families of random variables $\{Y_n\}$ and $\{Z_n\}$,

$$Y_n := \left(\frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_j \right)_{i=1}^N, \quad Z_n := \left(\frac{1}{n} \sum_{j=1}^{[n\psi_i^n]} U_j \right)_{i=1}^N,$$

which have laws μ_n and ν_n , respectively, as in (3.3)–(3.4). Since $\varphi^n \rightarrow \varphi$ we know that for any $\varepsilon > 0$ there exists an M_ε such that for all $n > M_\varepsilon$ we have that $\max_i |\varphi_i^n - \varphi_i| < \varepsilon/N$, and thus $|\psi_i^n - \psi_i| < \varepsilon$.

We have to show that for any $\delta > 0$

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\|Y_n - Z_n\|_\infty > \delta) = -\infty.$$

For $i \leq N$ consider the absolute difference between $Y_{n,i}$ and $Z_{n,i}$, i.e.,

$$(3.19) \quad |Y_{n,i} - Z_{n,i}| = \left| \frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_j - \frac{1}{n} \sum_{j=1}^{[n\psi_i^n]} U_j \right|.$$

Next we have that for any $n > M_\varepsilon$ it holds that $|n\psi_i^n - n\psi_i| < n\varepsilon$, which yields the upper bound,

$$|[n\psi_i^n] - [n\psi_i]| \leq [n\varepsilon] + 2,$$

since the rounded numbers differ at most by 1 from their real counterparts. This means that the difference of the two sums in (3.19) can be bounded by at most $[n\varepsilon] + 2$ elements of the U_j , which are for convenience denoted by U_j^* . Recalling that the U_j are nonnegative, we obtain

$$\left| \frac{1}{n} \sum_{j=1}^{[n\psi_i]} U_j - \frac{1}{n} \sum_{j=1}^{[n\psi_i^n]} U_j \right| \leq \frac{1}{n} \sum_{j=1}^{[n\varepsilon]+2} U_j^*.$$

Next we bound the probability that the difference exceeds δ , by using the above inequality:

$$\mathbb{P}(|Y_n - Z_n| > \delta) \leq \mathbb{P}\left(\frac{1}{n} \sum_{j=1}^{[n\varepsilon]+2} U_j^* > \delta\right) \leq \mathbb{E}[\exp(\theta U_1)]^{[n\varepsilon]+2} e^{-n\delta\theta},$$

where the last inequality follows from the Chernoff bound [DZ98, Eqn. (2.2.12)] for $\theta > 0$. Taking the log of this probability, dividing by n , and taking the limsup on both sides results in

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(|Y_n - Z_n|_\infty > \delta) \leq \varepsilon \Lambda_U(\theta) - \delta\theta.$$

By the assumption, $\Lambda_U(\theta) < \infty$ for all θ . Thus, $\varepsilon \rightarrow 0$ yields

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(|Y_n - Z_n|_\infty > \delta) \leq -\delta\theta.$$

As θ was arbitrary, the exponential equivalence follows by letting $\theta \rightarrow \infty$. \square

We conclude this section with some examples.

Example 3.7. Assume that the loss amounts have finite support, say on the interval $[0, u]$. Then we clearly have

$$\Lambda_U(\theta) = \log \mathbb{E}[e^{\theta U}] \leq \theta u < \infty.$$

So for any distribution with finite support, the assumption for Lemma 3.4 is satisfied, and thus Theorem 3.2 holds. Here, the i.i.d. default times, τ_i , can have an arbitrary discrete distribution on the time grid $\{1, \dots, N\}$.

In practical applications one (always) chooses a distribution with finite support for the loss amounts, since the exposure to every obligor is finite. Theorem 3.2 thus clearly holds for any (realistic) model of the loss given default.

An explicit expression for the rate function (3.5), or even the Fenchel-Legendre transform, is usually not available. On the other hand one can use numerical optimization techniques to calculate these quantities. \diamond

We next present an example to which Lemma 3.4 applies.

Example 3.8. Assume that the loss amount U is measured in a certain unit, and takes on the values $u, 2u, \dots$ for some $u > 0$. Assume that it has a distribution of Poisson type with parameter $\lambda > 0$, in the sense that for $i = 1, 2, \dots$

$$\mathbb{P}(U = (i+1)u) = e^{-\lambda} \frac{\lambda^i}{i!}.$$

It is then easy to check that $\Lambda_U(\theta) = \theta u + \lambda(e^{\theta u} - 1)$, being finite for all θ . Further calculations yield

$$\Lambda_U^*(x) = \left(\frac{x}{u} - 1\right) \log \left(\frac{1}{\lambda} \left(\frac{x}{u} - 1\right)\right) - \left(\frac{x}{u} - 1\right) + \lambda$$

for all $x > u$, and ∞ otherwise. Dividing this expression by x and letting $x \rightarrow \infty$, we observe that the resulting ratio goes to ∞ . As a consequence, Remark 3.6 now entails that Theorem 3.2 applies. It can also be argued that for any distribution U with tail behavior comparable to that of a Poisson distribution, Theorem 3.2 applies as well. \diamond

4. EXACT ASYMPTOTIC RESULTS

In the previous section we have established a sample-path large deviation principle on a finite time grid; this LDP provides us with logarithmic asymptotics of the probability that the sample path of $L_n(\cdot)/n$ is contained in a given set, say A . The results presented in this section are different in several ways. In the first place, we derive exact asymptotics (rather than logarithmic asymptotics). In the second place, our time domain is not assumed to be finite; instead we consider all integer numbers, \mathbb{N} . The price to be paid is that we restrict ourselves to special sets A , viz. those corresponding to the loss process (or the increment of the loss process) exceeding a given function. We work under the setup that we introduced in Section 2.1.

4.1. Crossing a Barrier. In this section we consider the asymptotic behavior of the probability that the loss process at some point in time is above a time-dependent level ζ . More precisely, we consider the set

$$(4.1) \quad A := \{f : T \rightarrow \mathbb{R}_0^+ \mid \exists t \in T : f(t) \geq \zeta(t)\},$$

for some function $\zeta(t)$ satisfying

$$(4.2) \quad \zeta(t) > \mathbb{E}[UZ(t)] = \mathbb{E}[U] F_t \quad \text{for all } t \in T,$$

with F_t as in (2.3). If we would consider a function ζ that does not satisfy (4.2), we are not in a large-deviations setting, in the sense that the probability of the event $\{L_n(\cdot)/n \in A\}$ converges to 1 by the law of large numbers. In order to obtain a more interesting result, we thus limit ourselves to levels that satisfy (4.2). For such levels we state the first main result of this section.

Theorem 4.1. *Assume that*

$$(4.3) \quad \text{there is a unique } t^* \in T \text{ such that } I_{UZ}(t^*) = \min_{t \in T} I_{UZ}(t),$$

and that

$$(4.4) \quad \liminf_{t \rightarrow \infty} \frac{I_{UZ}(t)}{\log t} > 0,$$

where $I_{UZ}(t) = \sup_{\theta} \{\theta \zeta(t) - \Lambda_{UZ(t)}(\theta)\} = \Lambda_{UZ(t)}^(\zeta(t))$. Then*

$$(4.5) \quad \mathbb{P}\left(\frac{1}{n}L_n(\cdot) \in A\right) = \frac{e^{-nI_{UZ}(t^*)}C^*}{\sqrt{n}} \left(1 + O\left(\frac{1}{n}\right)\right),$$

for A as in (4.1) and σ^ is such that $\Lambda'_{UZ(t^*)}(\sigma^*) = \zeta(t^*)$. The constant C^* follows from the Bahadur-Rao theorem (Theorem A.8), with $C^* = C_{UZ(t^*), \zeta(t^*)}$.*

Before proving our result, which will rely on arguments similar to those in [LM99], we first discuss the meaning and implications of Theorem 4.1. In addition we reflect on the role played by the assumptions. We do so by a sequence of remarks.

Remark 4.2. Comparing Theorem 4.1 to the Bahadur-Rao theorem (Theorem A.8), we observe that the probability of a sample mean exceeding a rare value has the same type of decay as the probability of our interest (i.e., the probability that the normalized loss process $L_n(\cdot)/n$ ever exceeds some function ζ). This decay looks like Ce^{-nI}/\sqrt{n} for positive constants C and I . This similarity can be explained as follows.

First observe that the probability of our interest is actually the probability of a *union* events. Evidently, this probability is larger than the probability of any of the events in this union, and hence also larger than the largest among these:

$$(4.6) \quad \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right) \geq \sup_{t \in T} \mathbb{P} \left(\frac{1}{n} L_n(t) \geq \zeta(t) \right).$$

Theorem 4.1 indicates that the inequality in (4.6) is actually tight (under the conditions stated). Informally, this means that the contribution of the maximizing t in the right-hand side of (4.6), say t^* , dominates the contributions of the other time epochs as n grows large. This essentially says that *given* that the rare event under consideration occurs, with overwhelming probability it happens at time t^* . \diamond

As is clear from the statement of Theorem 4.1, two assumptions are needed to prove the claim; we now briefly comment on the role played by these.

Remark 4.3. Assumption (4.3) is needed to make sure that there is not a time epoch \bar{t} , different from t^* , having a contribution of the same order as t^* . It can be verified from our proof that if the uniqueness assumption is not met, the probability under consideration remains asymptotically proportional to e^{-nI}/\sqrt{n} , but we lack a clean expression for the proportionality constant.

Assumption (4.4) has to be imposed to make sure that the contribution of the ‘upper tail’, that is, time epochs $t \in \{t^* + 1, t^* + 2, \dots\}$, can be neglected; more formally, we should have

$$\mathbb{P} \left(\exists t \in \{t^* + 1, t^* + 2, \dots\} : \frac{1}{n} L_n(t) \geq \zeta(t) \right) = o \left(\mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right) \right).$$

In order to achieve this, the probability that the normalized loss process exceeds ζ for large t should be sufficiently small. \diamond

Remark 4.4. We now comment on what Assumption (4.4) means. Clearly,

$$\Lambda_{UZ(t)}(\theta) = \log \mathbb{P}(\tau \leq t) \mathbb{E}[e^{\theta U}] + \mathbb{P}(\tau > t) \leq \log \mathbb{E}[e^{\theta U}],$$

as t grows, $\theta \geq 0$; the limiting value is actually $\log \mathbb{E}[e^{\theta U}]$ if τ is non-defective. This entails that

$$I_{UZ}(t) = \Lambda_{UZ(t^*)}^*(\zeta(t)) \geq \Lambda_U^*(\zeta(t)) = \sup_{\theta} (\theta \zeta(t) - \log \mathbb{E}[e^{\theta U}]).$$

We observe that Assumption (4.4) is fulfilled if $\liminf_{t \rightarrow \infty} \Lambda_U^*(\zeta(t))/\log t > 0$, which turns out to be valid under extremely mild conditions. Indeed, relying on Lemma A.4, we have that in great generality it holds that $\Lambda_U^*(x)/x \rightarrow \infty$ as $x \rightarrow \infty$. Then clearly any $\zeta(t)$, for which $\liminf_t \zeta(t)/\log t > 0$, satisfies Assumption (4.4), since

$$\liminf_{t \rightarrow \infty} \frac{\Lambda_U^*(\zeta(t))}{\log t} = \liminf_{t \rightarrow \infty} \frac{\Lambda_U^*(\zeta(t))}{\zeta(t)} \frac{\zeta(t)}{\log t}.$$

Alternatively, if U is chosen distributed exponentially with mean λ (which does not satisfy the conditions of Lemma A.4), then $\Lambda_U^*(t) = \lambda t - 1 - \log(\lambda t)$, such that we have that

$$\liminf_{t \rightarrow \infty} \frac{I_U(\log t)}{\log t} = \lambda > 0.$$

Barrier functions ζ that grow at a rate slower than $\log t$, such as $\log \log t$, are in this setting clearly not allowed. \diamond

Proof of Theorem 4.1. We start by rewriting the probability of interest as

$$\mathbb{P}\left(\frac{1}{n}L_n(\cdot) \in A\right) = \mathbb{P}\left(\exists t \in T : \frac{L_n(t)}{n} \geq \zeta(t)\right).$$

For an arbitrary point k in T we have

$$(4.7) \quad \mathbb{P}\left(\exists t \in T : \frac{L_n(t)}{n} \geq \zeta(t)\right) \leq \mathbb{P}\left(\exists t \leq k : \frac{L_n(t)}{n} \geq \zeta(t)\right) + \mathbb{P}\left(\exists t > k : \frac{L_n(t)}{n} \geq \zeta(t)\right).$$

We first focus on the second part in (4.7). We can bound this by

$$\begin{aligned} \mathbb{P}\left(\exists t > k : \frac{L_n(t)}{n} \geq \zeta(t)\right) &\leq \sum_{i=k+1}^{\infty} \mathbb{P}\left(\frac{L_n(i)}{n} \geq \zeta(i)\right) \\ &\leq \sum_{i=k+1}^{\infty} \inf_{\theta > 0} \mathbb{E}\left[\exp\left(\theta \sum_{j=1}^n U_j Z_j(i)\right)\right] e^{-n\zeta(i)\theta}, \end{aligned}$$

where the second inequality is due to the Chernoff bound [DZ98, Eqn. (2.2.12)]. The independence between the U_i and $Z_i(t)$, together with the assumption that the U_i are i.i.d. and the $Z_i(t)$ are i.i.d., yields

$$\begin{aligned} \sum_{i=k+1}^{\infty} \inf_{\theta > 0} \mathbb{E}\left[\exp\left(\theta \sum_{j=1}^n U_j Z_j(i)\right)\right] e^{-n\zeta(i)\theta} &= \sum_{i=k+1}^{\infty} \inf_{\theta > 0} \prod_{j=1}^n \mathbb{E}[\exp(\theta U_j Z_j(i))] e^{-n\zeta(i)\theta} \\ &= \sum_{i=k+1}^{\infty} \exp\left(-n \sup_{\theta > 0} (\zeta(i)\theta - \Lambda_{UZ(i)}(\theta))\right) = \sum_{i=k+1}^{\infty} \exp(-nI_{UZ}(\zeta(i))). \end{aligned}$$

By (4.4) we have that

$$\liminf_{t \rightarrow \infty} \frac{I_{UZ}(t)}{\log t} = \beta,$$

for some $\beta > 0$ (possibly ∞). Hence there exists an m such that for all $i > m$

$$(4.8) \quad I_{UZ}(i) > \alpha \log i > I_{UZ}(t^*),$$

where $\alpha = \beta/2$ (in case $\beta = \infty$, any $0 < \alpha < \infty$ suffices) and t^* defined in (4.3). Choosing $k = m$, we obtain by using the first inequality in (4.8) for $n > 1/\alpha$

$$\sum_{i=m+1}^{\infty} \exp(-nI_{UZ}(\zeta(i))) \leq \sum_{i=m+1}^{\infty} \exp(-n\alpha \log i) \leq \frac{1}{n\alpha - 1} \exp((-n\alpha + 1) \log m),$$

where the last inequality trivially follows by bounding the summation (from above) by an appropriate integral. Next we multiply and divide this by $\mathbb{P}(L_n(t^*)/n > \zeta(t^*))$ and we apply the Bahadur-Rao theorem, which results in

$$\begin{aligned} \frac{1}{n\alpha - 1} e^{(-n\alpha + 1) \log m} &= \frac{1}{n\alpha - 1} e^{(-n\alpha + 1) \log m} \frac{\mathbb{P}(L_n(t^*)/n > \zeta(t^*))}{\mathbb{P}(L_n(t^*)/n > \zeta(t^*))} \\ &= \mathbb{P}\left(\frac{1}{n}L_n(t^*) > \zeta(t^*)\right) \frac{m\sqrt{n}C^*}{n\alpha - 1} \left(1 + O\left(\frac{1}{n}\right)\right) e^{-n(\alpha \log m - I_{UZ}(t^*))}. \end{aligned}$$

The second inequality in (4.8) yields $\alpha \log m - I_{UZ}(t^*) > \delta$, for some $\delta > 0$. Applying this inequality, we see that this bounds the second term in (4.7), in the sense that as $n \rightarrow \infty$,

$$\mathbb{P}\left(\exists t > k : \frac{L_n(t)}{n} \geq \zeta(t)\right) / \mathbb{P}\left(\frac{1}{n}L_n(t^*) > \zeta(t^*)\right) \rightarrow 0.$$

To complete the proof we need to bound the first term of (4.7), where we use that $k = m$. For this we again use the Bahadur-Rao theorem. Next to this theorem we use the uniqueness of t^* , which implies that for $i \leq m$ and $i \neq t^*$ there exists an $\varepsilon^* > 0$, such that

$$I_{UZ}(t^*) + \varepsilon^* \leq I_{UZ}(i).$$

This observation yields, with σ_i such that $\Lambda'_{UZ(i)}(\sigma_i) = \zeta(i)$,

$$\begin{aligned} \mathbb{P}\left(\exists t \leq m : \frac{L_n(t)}{n} \geq \zeta(t)\right) &\leq \sum_{i=1}^m \mathbb{P}\left(\frac{L_n(i)}{n} \geq \zeta(i)\right) \\ &\leq \mathbb{P}\left(\frac{1}{n}L_n(t^*) > \zeta(t^*)\right) \left(1 + O\left(\frac{1}{n}\right)\right) \left(\sum_{i=1}^m \frac{C^*}{C_{UZ(i),\zeta(i)}} \frac{e^{-nI_{UZ}(t_i)}}{e^{-nI_{UZ}(t^*)}}\right) \\ &\leq \mathbb{P}\left(\frac{1}{n}L_n(t^*) > \zeta(t^*)\right) \left(1 + O\left(\frac{1}{n}\right)\right) \left(1 + m \times \max_{i=1,\dots,m} \left(\frac{C^*}{C_{UZ(i),\zeta(i)}}\right) e^{-n\varepsilon^*}\right) \\ &= \mathbb{P}\left(\frac{1}{n}L_n(t^*) > \zeta(t^*)\right) \left(1 + O\left(\frac{1}{n}\right)\right) \left(1 + O\left(e^{-n\varepsilon^*}\right)\right) \end{aligned}$$

Combining the above findings, we observe

$$\mathbb{P}\left(\exists t \in T : \frac{L_n(t)}{n} \geq \zeta(t)\right) \leq \mathbb{P}\left(\frac{L_n(t^*)}{n} \geq \zeta(t^*)\right) \left(1 + O\left(\frac{1}{n}\right)\right).$$

Together with the trivial bound

$$\mathbb{P}\left(\exists t \in T : \frac{L_n(t)}{n} \geq \zeta(t)\right) \geq \mathbb{P}\left(\frac{L_n(t^*)}{n} \geq \zeta(t^*)\right),$$

this yields

$$\mathbb{P}\left(\exists t \in T : \frac{L_n(t)}{n} \geq \zeta(t)\right) = \mathbb{P}\left(\frac{L_n(t^*)}{n} > \zeta(t^*)\right) \left(1 + O\left(\frac{1}{n}\right)\right).$$

Applying the Bahadur-Rao theorem to the right-hand side of the previous display yields the desired result. \square

4.2. Large Increments of the Loss Process. In the previous section we identified the asymptotic behavior of the probability that at some point in time the normalized loss process $L_n(\cdot)/n$ exceeds a certain level. We can carry out a similar procedure to obtain insight in the large deviations of the *increments* of the loss process. Here we consider times where the increment of the loss between time s and t exceeds a threshold $\xi(s, t)$. More precisely, we consider the event

$$(4.9) \quad A := \{f : T \times T \rightarrow \mathbb{R}_0^+ \mid \exists s < t \in T : f(s, t) \geq \xi(s, t)\}.$$

Being able to deal with events of this type, we can for instance analyze the likelihood of the occurrence of a large loss during a short period; we remark that with the event (4.1) from the previous subsection, one cannot distinguish the cases where the loss is zero for all times before t and $x > \zeta(t)$ at time t , and the case where the loss is just below the level ζ for all times before time t and then ends up at x at time t . Clearly, events of the (4.9) make it possible to distinguish between such paths.

In order to avoid trivial results, we impose a condition similar to (4.2), namely

$$(4.10) \quad \xi(s, t) > \mathbb{E}[U] (F_t - F_s),$$

for all $s < t$. The law of large numbers entails that for functions ξ that do not satisfy this condition, the probability under consideration does not correspond to a rare event.

A similar probability has been considered in [DDD04], where the authors derive the logarithmic asymptotic behavior of the probability that the increment of the loss, for some $s < t$, in a bounded interval exceeds a thresholds that depends only on $t - s$. In contrast, our approach uses a more flexible threshold, which depends on both times s and t , and in addition we derive the *exact* asymptotic behavior of this probability.

Theorem 4.5. *Assume that*

$$(4.11) \quad \text{there is a unique } s^* < t^* \in T \text{ such that } I_{UZ}(s^*, t^*) = \min_{s < t} I_{UZ}(s, t),$$

and that

$$(4.12) \quad \inf_{s \in T} \liminf_{t \rightarrow \infty} \frac{I_{UZ}(s, t)}{\log t} > 0,$$

where $I_{UZ}(s, t) = \sup_{\theta} (\theta \xi(s, t) - \Lambda_{U(Z(t)-Z(s))}(\theta)) = \Lambda_{U(Z(t)-Z(s))}^(\xi(s, t))$. Then*

$$(4.13) \quad \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \right) = \frac{e^{-n I_{UZ}(s^*, t^*)} C^*}{\sqrt{n}} \left(1 + O \left(\frac{1}{n} \right) \right),$$

for A as in (4.10) and σ^ is such that $\Lambda'_{U(Z(t^*)-Z(s^*))}(\sigma^*) = \xi(s^*, t^*)$. The constant C^* follows from the Bahadur-Rao theorem (Theorem A.8), with $C^* = C_{U(Z(t^*)-Z(s^*))}(\xi(s^*, t^*))$.*

Remark 4.6. A first glance at Theorem 4.5 tells us the obtained result is very similar to the result of Theorem 4.1. The second condition, i.e., Inequality (4.12), however, seems to be more restrictive than the corresponding condition, i.e., Inequality (4.4), due to the infimum over s . This assumption has to make sure that the ‘upper tail’ is negligible for any s . In the previous subsection we have seen that, under mild restrictions, the upper tail can be safely ignored when the barrier function grows at a rate of at least $\log t$. We can extend this claim to our new setting of large increments, as follows.

First note that

$$\inf_{s \in T} \liminf_{t \rightarrow \infty} \frac{I_{UZ}(s, t)}{\log t} \geq \inf_{s \in T} \liminf_{t \rightarrow \infty} \frac{\Lambda_U^*(\xi(s, t))}{\log t}.$$

Then consider thresholds that, next to condition (4.10), satisfy that for all s

$$(4.14) \quad \liminf_{t \rightarrow \infty} \frac{\xi(s, t)}{\log t} > 0.$$

Then, under the conditions of Lemma A.4, we have that

$$(4.15) \quad \liminf_{t \rightarrow \infty} \frac{\Lambda_U^*(\xi(s, t))}{\log t} = \liminf_{t \rightarrow \infty} \frac{\Lambda_U^*(\xi(s, t))}{\xi(s, t)} \frac{\xi(s, t)}{\log t} = \infty,$$

since the second term remains positive by (4.14) and the first term tends to infinity by Lemma A.4. Having established (4.15) for all s , it is clear that (4.12) is satisfied.

The sufficient condition (4.14) shows that the range of admissible barrier functions is quite substantial, and, importantly, imposing (4.12) is not as restrictive as it seems at first glance. \diamond

Proof of Theorem 4.5. The proof of this theorem is very similar to that of Theorem 4.1. Therefore we only sketch the proof here.

As before, the probability of interest is split up into a ‘front part’ and ‘tail part’. The tail part can be bounded using Assumption (4.12); this is done analogously to the way Assumption (4.4) was used in the proof of Theorem 4.1. The uniqueness assumption (4.11) then shows that the probability of interest is asymptotically equal to the probability that the increment between time

s^* and t^* exceeds $\xi(s^*, t^*)$; this is an application of the Bahadur-Rao theorem. Another application of the Bahadur-Rao theorem to the probability that the increment between time s^* and t^* exceeds $\xi(s^*, t^*)$ yields the result. \square

5. DISCUSSION AND CONCLUDING REMARKS

In this paper we have established a number of results with respect to the asymptotic behavior of the distribution of the loss process. In this section we discuss some of the assumptions in more detail and we consider extensions of the results that we have derived.

5.1. Extensions of the Sample-Path LDP. The first part of our work, Section 3, was devoted to establishing a sample-path large deviation principle on a finite time grid. Here we modelled the loss process as the sum of i.i.d. loss amounts multiplied by i.i.d. default indicators. From a practical point of view one can argue that the assumptions underlying our model are not always realistic. In particular, the random properties of the obligors cannot always be assumed independent. In addition, the assumption that all obligors behave in an i.i.d. fashion will not necessarily hold in practice. Both shortcomings can be dealt with, however, by adapting the model slightly.

A common way to introduce dependence, taken from [DDD04], is by supposing that there is a ‘macro-environmental’ variable Y to which all obligors react, but conditional on which the loss epochs and loss amounts are independent. First observe that our results are then valid for any specific realization y of Y . Denoting the exponential decay rate by r_y , i.e.,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{1}{n} L_n(\cdot) \in A \mid Y = y \right) = r_y,$$

the *unconditional* decay rate is just the maximum over the r_y ; this is trivial to prove if Y can attain values in a finite set only. A detailed treatment of this is beyond the scope of this paper.

The assumption that all obligors have the same distribution can be relaxed to the case where we assume that there are m different classes of obligors (for instance determined by their default ratings). We further assume that each class i makes up a fraction a_i of the entire portfolio. Then we can extend the LDP of Theorem 3.2 to a more general one, by splitting up the loss process into m loss processes, each corresponding to a class. Conditioning on the realizations of these processes, we can derive the following rate function:

$$(5.1) \quad I_{U,p,m}(x) := \inf_{\varphi \in \Phi^m} \inf_{v \in V_x} \sum_{j=1}^m \sum_{i=1}^N a_i \varphi_i^j \left(\log \left(\frac{\varphi_i^j}{p_i^j} \right) + \Lambda_U^* \left(\frac{v_i^j}{a_i \varphi_i^j} \right) \right),$$

where $V_x = \{v \in \mathbb{R}_+^{m \times N} \mid \sum_{j=1}^m v_i^j = \Delta x_i \text{ for all } i \leq N\}$, and Φ^m is the Cartesian product $\Phi \times \dots \times \Phi$ (m times), with Φ as in (3.1). The optimization over the set V_x follows directly from conditioning on the realizations of the per-class loss processes. We leave out the formal derivation of this result; this multi-class case is notationally considerably more involved than the single-class case, but essentially all steps carry over.

In our sample-path LDP we assumed that defaults can only occur on a finite grid. While this assumption is justifiable from a practical point of view, an interesting mathematical question is whether it can be relaxed. In self-evident notation, one would expect that the rate function

$$I_{U,p,\infty}(x) := \inf_{\varphi \in \Phi_\infty} \sum_{i=1}^\infty \varphi_i \left(\log \left(\frac{\varphi_i}{p_i} \right) + \Lambda_U^* \left(\frac{\Delta x_i}{\varphi_i} \right) \right).$$

It can be checked, however, that the argumentation used in the proof of Theorem 3.2 does not work; in particular, the choice of a suitable topology plays an important role.

If losses can occur on a continuous entire interval, i.e., $[0, N]$, we expect, for a nondecreasing and differentiable path x , the rate function

$$(5.2) \quad I_{U,p,[0,N]}(x) := \inf_{\varphi \in \mathcal{M}} \int_0^N \varphi(t) \left(\log \left(\frac{\varphi(t)}{p(t)} \right) + \Lambda_U^* \left(\frac{x'(t)}{\varphi(t)} \right) \right) dt,$$

where \mathcal{M} is the space of all densities on $[0, N]$ and p the density of the default time τ . One can easily guess the validity of (5.2) from (3.5) by using Riemann sums to approximate the integral. A formal proof, however, requires techniques that are essentially different from the ones used to establish Theorem 3.2, and therefore we leave this for future research.

5.2. Extensions of the Exact Asymptotics. In the second part of the paper, i.e., Section 4, we have derived the exact asymptotic behavior for two special events. First we showed that, under certain conditions, the probability that the loss process exceeds a certain time-dependent level, is asymptotically equal to the probability that the process exceeds this level at the ‘most likely’ time t^* . The exact asymptotics of this probability are obtained by applying the Bahadur-Rao theorem. A similar result has been obtained for an event related to the *increment* of the loss process. One could think of refining the logarithmic asymptotics, as developed in Section 3, to exact asymptotics. Note, however, that this is far from straightforward, as for general sets these asymptotics do not necessarily coincide with those of a univariate random variable, cf. [MMNvU06].

REFERENCES

- [Aco94] A. de Acosta, *Large deviations for vector-valued Lévy processes*, Stochastic Processes and their Applications **51** (1994), 75–115.
- [ADEH99] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, *Coherent measures of risk*, Mathematical Finance **9** (1999), 203–228.
- [Aze80] R. Azencott, *Ecole d’Été de probabilités de Saint-Flour VIII-1978*, Lecture Notes in Mathematics, vol. 774, ch. Grandes déviations et applications, pp. 1–176, Springer, Berlin, 1980.
- [BR60] R.R. Bahadur and R. Ranga Rao, *On deviations of the sample mean*, Annals of Mathematical Statistics **31** (1960), 1015–1027.
- [Cha95] C.-S. Chang, *Sample path large deviations andintree networks*, Queueing Systems **20** (1995), 7–36.
- [DDD04] A. Dembo, J.-D. Deuschel, and D. Duffie, *Large portfolio losses*, Finance and Stochastics **8** (2004), 3–16.
- [DS03] D. Duffie and K. Singleton, *Credit risk: Pricing, measurement, and management*, Princeton University Press, 2003.
- [DZ98] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, 2nd ed., Springer Verlag, New York, 1998.
- [EKM97] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling extremal events*, Applications of Mathematics, vol. 33, Springer-Verlag, 1997.
- [GKS07] P. Glasserman, W. Kang, and P. Shahabuddin, *Large deviations in multifactor portfolio credit risk*, Mathematical Finance **17** (2007), 345–379.
- [Gla04] P. Glasserman, *Monte Carlo methods in financial engineering*, Applications of Mathematics, vol. 53, Springer-Verlag, 2004.
- [Gor02] M.B. Gordy, *Saddlepoint approximation of CreditRisk+*, Journal of Banking and Finance **26** (2002), 1337–1355.
- [Jor06] P. Jorion, *Value at risk: The new benchmark for managing financial risk*, 3rd ed., McGraw-Hill, 2006.
- [Kal86] P. Kall, *Approximation to optimization problems: An elementary review*, Mathematics of Operations Research **11** (1986), 9–18.
- [LKSS01] A. Lucas, P. Klaassen, P. Spreij, and S. Straetmans, *An analytical approach to credit risk in large corporate bond and loan portfolios*, Journal of Banking and Finance **25** (2001), 1635–1664.

- [LM99] N. Likhanov and R. Mazumdar, *Cell loss asymptotics for buffers fed with a large number of independent stationary sources*, Journal of Applied Probability **36** (1999), 86–96.
- [MFE05] A. McNeil, R. Frey, and P. Embrechts, *Quantitative risk management: Concepts, techniques and tools*, Princeton University Press, 2005.
- [MMNvU06] M. Mandjes, P. Mannersalo, I. Norros, and M. van Uitert, *Large deviations of infinite intersections of events in Gaussian processes*, Stochastic Processes and their Applications **116** (2006), 1269–1293.
- [Mog76] A. Mogul'skiĭ, *Large deviations for trajectories of multi dimensional random walks*, Theory of Probability and its Applications **21** (1976), 300–315.
- [MTB01] R. Martin, K. Thompson, and C. Browne, *Taking to the saddle*, Risk **14** (2001), 91–94.

APPENDIX A. BACKGROUND RESULTS

In this section we state a number of definitions and results, taken from [DZ98], which are used in the proofs in this paper.

Theorem A.1 (Cramér). *Let X_i be i.i.d. real valued random variables with all exponential moments finite and let μ_n be the law of the average $S_n = \sum_{i=1}^n X_i/n$. Then the sequence $\{\mu_n\}$ satisfies an LDP with rate function $\Lambda^*(\cdot)$, where Λ^* is the Fenchel-Legendre transform of the X_i .*

Proof. See for example [DZ98, Thm. 2.2.3]. □

Definition A.2. *We say that two families of measures $\{\mu_n\}$ and $\{\nu_n\}$ on a metric space (\mathcal{X}, d) are exponentially equivalent if there exist two families of \mathcal{X} -valued random variables $\{Y_n\}$ and $\{Z_n\}$ with marginal distributions $\{\mu_n\}$ and $\{\nu_n\}$, respectively, such that for all $\delta > 0$*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(d(X_n, Y_n) \geq \delta) = -\infty.$$

Lemma A.3. *For every triangular array $a_n^i \geq 0$, $n \geq 1$, $1 \leq i \leq n$,*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \sum_{i=1}^n a_n^i = \limsup_{n \rightarrow \infty} \max_{i=1, \dots, n} \frac{1}{n} \log a_n^i.$$

Proof. Elementary, but also a direct consequence of [DZ98, Lemma 1.2.15]. □

Lemma A.4. *Let $\Lambda(\theta) < \infty$ for all $\theta \in \mathbb{R}$, then*

$$\lim_{|x| \rightarrow \infty} \frac{\Lambda^*(x)}{|x|} = \infty.$$

Proof. This result is a part of [DZ98, Lemma 2.2.20]. □

Lemma A.5. *Let $K_{n,i}$ be defined as $K_{n,j} := \#\{i \in \{1, \dots, n\} \mid \tau_i = j\}$. Then for any vector $k \in \mathbb{N}^N$, such that $\sum_{i=1}^N k_i = n$, we have that*

$$(n+1)^{-N} \exp(-nH(k \mid p)) \leq \mathbb{P}(K_n = k) \leq \exp(-nH(k \mid p)),$$

where

$$H(k \mid p) = \sum_{i=1}^N \frac{k_i}{n} \log \left(\frac{k_i}{np_i} \right),$$

and p_i as defined in (2.2).

Proof. See [DZ98, Lemma 2.1.9]. □

Lemma A.6. *Define*

$$Z_n(t) := \frac{1}{n} \sum_{i=1}^{\lfloor nt \rfloor} X_i, \quad 0 \leq t \leq 1,$$

for an i.i.d. sequence of \mathbb{R}^d valued random variables X_i . Let μ_n denote the law of $Z_n(\cdot)$ in $L_\infty([0, 1])$. For any discretization $J = \{0 < t_1 < \dots < t_{|J|} \leq 1\}$ and any $f : [0, 1] \rightarrow \mathbb{R}^d$, let $p_J(f)$ denote the vector $(f(t_i))_{i=1}^{|J|} \in (\mathbb{R}^d)^{|J|}$. Then the sequence of laws $\{\mu_n \circ p_J^{-1}\}$ satisfies the LDP in $(\mathbb{R}^d)^{|J|}$ with the good rate function

$$I_J(z) = \sum_{i=1}^{|J|} (t_i - t_{i-1}) \Lambda^* \left(\frac{z_i - z_{i-1}}{t_i - t_{i-1}} \right),$$

where Λ^* is the Fenchel-Legendre transform of X_1 .

Proof. See [DZ98, Lemma 5.1.8]. This lemma is one of the key steps in proving Mogul'skiĭ's theorem, which provides a sample-path LDP for $Z_n(\cdot)$ on a bounded interval. \square

Theorem A.7. *If an LDP with a good rate function $I(\cdot)$ holds for the probability measures $\{\mu_n\}$, which are exponentially equivalent to $\{\nu_n\}$, then the same LDP holds for $\{\nu_n\}$.*

Proof. See [DZ98, Thm. 4.2.13]. \square

Theorem A.8 (Bahadur-Rao). *Let X_i be an i.i.d. real-valued sequence random variables. Then we have*

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n X_i \geq q \right) = \frac{e^{-n\Lambda_X^*(q)} C_{X,q}}{\sqrt{n}} \left(1 + O \left(\frac{1}{n} \right) \right).$$

The constant $C_{X,q}$ depends on the type of distribution of X_1 , as specified by the following two cases.

- (i) *The law of X_1 is lattice, i.e. for some x_0, d , the random variable $(X_1 - x_0)/d$ is (a.s.) an integer number, and d is the largest number with this property. Under the additional condition $0 < \mathbb{P}(X_1 = q) < 1$, the constant $C_{X,q}$ is given by*

$$C_{X,q} = \frac{d}{(1 - e^{-\sigma d}) \sqrt{2\pi \Lambda_X''(\sigma)}},$$

where σ satisfies $\Lambda_X'(\sigma) = q$.

- (ii) *If the law of X_1 is non-lattice, the constant $C_{X,q}$ is given by*

$$C_{X,q} = \frac{1}{\sigma \sqrt{2\pi \Lambda_X''(\sigma)}},$$

with σ as in case (i).

Proof. We refer to [BR60] or [DZ98, Thm. 3.7.4] for the proof of this result. \square

Acknowledgment. VL would like to thank ABN AMRO bank for providing financial support. Part of this work was carried out while MM was at Stanford University, US. The authors are indebted to E.J. Balder (Utrecht University, the Netherlands) for pointing out to us the relevance of epi-convergence to our research.

KORTEWEG-DE VRIES INSTITUTE FOR MATHEMATICS, UNIVERSITY OF AMSTERDAM, THE NETHERLANDS — ABN AMRO, AMSTERDAM, THE NETHERLANDS

E-mail address: v.j.g.leijdekker@uva.nl

KORTEWEG-DE VRIES INSTITUTE FOR MATHEMATICS, UNIVERSITY OF AMSTERDAM, THE NETHERLANDS — EURANDOM, EINDHOVEN, THE NETHERLANDS — CWI, AMSTERDAM, THE NETHERLANDS

E-mail address: m.r.h.mandjes@uva.nl

KORTEWEG-DE VRIES INSTITUTE FOR MATHEMATICS, UNIVERSITY OF AMSTERDAM, THE NETHERLANDS

E-mail address: p.j.c.spreij@uva.nl